

Spatially adaptive covariance tapering

DAVID BOLIN AND JONAS WALLIN

Mathematical Sciences, Chalmers and University of Gothenburg

Abstract: Covariance tapering is a popular approach for reducing the computational cost of spatial prediction and parameter estimation for Gaussian process models. However, tapering can have poor performance when the process is sampled at spatially irregular locations or when non-stationary covariance models are used. This work introduces an adaptive tapering method in order to improve the performance of tapering in these problematic cases. This is achieved by introducing a computationally convenient class of compactly supported non-stationary covariance functions, combined with a new method for choosing spatially varying taper ranges. Numerical experiments are used to show that the performance of both kriging prediction and parameter estimation can be improved by allowing for spatially varying taper ranges. However, although adaptive tapering outperforms regular tapering, simply dividing the data into blocks and ignoring the dependence between the blocks is often a better method for parameter estimation.

Key words: Kriging; Sparse matrices; Compactly supported covariances; non-stationary covariances; Maximum likelihood

1 Introduction

Gaussian processes are important for statistical analysis of spatial data. In applications the goal is often to predict the process at unobserved locations, which is done by computing the conditional mean of the field given the data. This is referred to as the kriging prediction in geostatistics, and requires solving $\Sigma^{-1}\mathbf{x}$ where \mathbf{x} is a vector with observed values of the field and Σ is the covariance matrix for the field at the observation locations. Thus, if the field is observed at N locations, the computational cost for kriging prediction is in general $\mathcal{O}(N^3)$. This limits the applicability for large datasets and is commonly referred to as the “big N” problem.

Covariance tapering is a popular method for handling the big N problem. The basic idea of this method is to set small elements in Σ to zero, which enables the use of computationally efficient sparse matrix techniques for computing $\Sigma^{-1}\mathbf{x}$. For spatial problems, this typically reduces the computational complexity from $\mathcal{O}(N^3)$ to $\mathcal{O}(N^{3/2})$. The simplest way to introduce zeros in Σ is to replace the covariance function, $r(\mathbf{h})$, of the Gaussian process with some compactly supported covariance function $T(\mathbf{h})$, such as a Wendland function (Wendland, 1995; Gneiting, 2002). An alternative is to replace $r(\mathbf{h})$ with a tapered version $r_{tap}(\mathbf{h}) = r(\mathbf{h})T(\mathbf{h})$. The covariance range of $T(\mathbf{h})$, referred to as the taper range, determines the sparsity of the resulting covariance matrix.

Several authors have studied the effect of tapering in the case when $r(\mathbf{h})$ is a Matérn covariance function (Matérn, 1960). Kaufman et al. (2008) showed that certain parameters of the Matérn covariance function can be consistently estimated using tapering, and Du et al. (2009), Shaby and Ruppert (2012), and Wang et al. (2011) further studied the asymptotic properties of tapered estimators in the case of Matérn covariance models. Furrer et al. (2006) showed that tapering can be used with asymptotically negligible loss in the case of kriging

prediction for Matérn fields, and Stein (2013) extended some of these results beyond the Matérn model.

However, asymptotic results do not guarantee that the method is practically useful. Stein (2013) used numerical experiments to show that covariance tapering often does not work as well for parameter estimation as simply splitting the observations into blocks and ignoring the dependence between the blocks. Furthermore, Bolin and Lindgren (2013) showed that tapering also can perform poorly for kriging prediction compared with other methods for handling the big N problem. A reason for this is that the computational cost for stationary tapering depends on the ratio of the taper range and the average spacing of the observations, whereas the accuracy depends on the ratio of the taper range and the range of $r(\mathbf{h})$. The accuracy also depends on how the observations are located spatially, and more sparsely observed areas will have higher errors for tapered kriging predictions. This often leads to higher errors close to the boundaries of the observation domain as well as “spotty” predicted surfaces where sparsely observed regions are biased towards zero.

Stein (2013) proposed using a simple deformation method to improve performance near the boundaries of the domain. However, this method only works for rectangular domains and does not help in sparsely observed regions. Anderes et al. (2013) proposed a more sophisticated deformation method based on quasi-conformal maps where an adaptive taper $T(\mathbf{s}, \mathbf{t}) = T(\|\varphi(\mathbf{s}) - \varphi(\mathbf{t})\|)$ was defined using a stationary taper T and a warping function φ . A drawback with this approach is that estimation of the warping function φ is highly computationally demanding, which limits the usefulness of the method.

This work introduces a new adaptive taper method that is more flexible than that by Stein (2013) while at the same time being more computationally efficient than the method by Anderes et al. (2013). This is achieved by first introducing a computationally convenient class of compactly supported non-stationary covariance functions in Section 2. How to use this for kriging prediction is then discussed in Section 3. The section first introduces a new method for choosing spatially varying taper ranges and then investigates the accuracy of the method using simulated data. Section 4 presents a method for using the adaptive tapers for parameter estimation and investigates the accuracy of the method using numerical experiments. The conclusion of these comparisons is that adaptive tapering outperforms regular tapering, but simply dividing the data into blocks and ignoring the dependence between the blocks is often a better method for parameter estimation. This is in line with what Stein (2013) found. Finally, Section 5 contains comments and suggestions for future research.

2 Adaptive covariance tapers

A popular method for constructing non-stationary covariance models is to use the process convolution approach (Barry and Ver Hoef, 1996; Higdon, 2001; Cressie and Ravlicová, 2002; Rodrigues and Diggle, 2010), where a Gaussian stochastic field, $X(\mathbf{s})$, is defined as the convolution of a Brownian sheet \mathcal{B} and some convolution kernel $k_s(\mathbf{u})$. The covariance function of $X(\mathbf{s})$ is then given by the integral $C(\mathbf{s}, \mathbf{t}) = \int k_s(\mathbf{u})k_t(\mathbf{u}) d\mathbf{u}$, which is non-stationary if the kernel changes spatially. One can also note that the support of the covariance function is determined by the support of the kernel. The idea here is therefore to use compactly supported and non-stationary kernels in the process-convolution approach to construct adaptive tapering functions. Thus, we choose the tapering function $T(\mathbf{s}, \mathbf{t})$ as

$$T(\mathbf{s}, \mathbf{t}) = \int k_s(\mathbf{u})k_t(\mathbf{u}) d\mathbf{u}. \quad (1)$$

The advantage with this approach is that $T(\mathbf{s}, \mathbf{t})$ by construction is a valid covariance function for any square integrable $k_s(\mathbf{u})$. The disadvantage, however, is that numerical methods

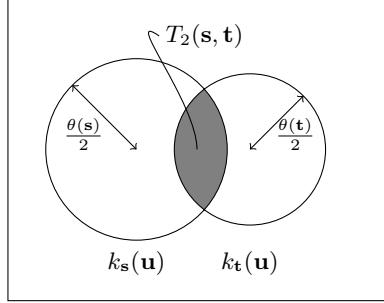


Figure 1: The value of the non-stationary covariance function $T_2(\mathbf{s}, \mathbf{t})$ is given by the area of the shaded region normalized by the areas of the two circles.

often are required to evaluate the integral in (1). This is computationally expensive and therefore reduces usefulness of the construction. A crucial property of the tapers we construct below is that the integral can be solved analytically, which means that one has an explicit form of the non-stationary taper function.

2.1 Hyperspherical tapers

The perhaps simplest choice of $k_s(\mathbf{u})$ is the indicator function

$$k_{2,s}(\mathbf{u}) = \frac{1}{\theta(s)\sqrt{\pi}} \mathbb{I}\left(\|\mathbf{s} - \mathbf{u}\| < \frac{\theta(s)}{2}\right) = \begin{cases} \frac{1}{\theta(s)\sqrt{\pi}} & \text{if } \|\mathbf{s} - \mathbf{u}\| < \theta(s)/2, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

The normalization of the kernel is chosen so that the resulting taper satisfies $T(\mathbf{s}, \mathbf{s}) = 1$. In this parameterization, the taper range of a stationary model with $\theta(\mathbf{s}) \equiv \theta$ is given by θ . The value of $T(\mathbf{s}, \mathbf{t})$ is in this case obtained as the normalized area of the asymmetric lens produced by the intersection of two circles centered at \mathbf{s} and \mathbf{t} , with diameters $\theta(\mathbf{s})$ and $\theta(\mathbf{t})$ respectively. See Figure 1 for an illustration.

Straightforward calculations give that the resulting taper is

$$T_2(\mathbf{s}, \mathbf{t}) = \frac{1}{\pi r_{st} R_{st}} \cdot \begin{cases} \pi r_{st} & d_{st} < R_{st} - r_{st}, \\ V_2\left(R_{st}, \frac{d_{st}^2 + R_{st}^2 - r_{st}^2}{2d_{st}}\right) + V_2\left(r_{st}, \frac{d_{st}^2 + r_{st}^2 - R_{st}^2}{2d_{st}}\right) & R_{st} - r_{st} \leq d_{st} < R_{st} + r_{st}, \\ 0 & \text{otherwise.} \end{cases}$$

Here $d_{st} = \|\mathbf{s} - \mathbf{t}\|$, $R_{st} = \frac{1}{2} \max(\theta(\mathbf{s}), \theta(\mathbf{t}))$, and $r_{st} = \frac{1}{2} \min(\theta(\mathbf{s}), \theta(\mathbf{t}))$. Further,

$$V_2(r, x) = \begin{cases} r^2 \cos^{-1}\left(\frac{x}{r}\right) - x\sqrt{r^2 - x^2} & |x| < r \\ 0 & \text{otherwise,} \end{cases} \quad (3)$$

is the area of a circular cap with triangular height x of a circle with radius r . The case $d_{st} < R_{st} - r_{st}$ does not occur if the taper range $\theta(\mathbf{s})$ is Hölder continuous with exponent 1. This can be a natural condition to require in applications where the taper range should vary smoothly across space.

With $\theta(\mathbf{s}) \equiv \theta$, the taper function reduces to

$$T(d) = \begin{cases} \frac{2}{\pi} \cos^{-1}\left(\frac{d}{\theta}\right) - \frac{2}{\pi} \frac{d}{\theta} \sqrt{1 - \frac{d^2}{\theta^2}} & \text{for } d < \theta, \\ 0 & \text{otherwise,} \end{cases}$$

and it is easy to show that this is a valid covariance function for processes on \mathbb{R} and on \mathbb{R}^2 , but not for any \mathbb{R}^m with $m > 2$.

A variant of the taper is obtained if \mathbf{s} and \mathbf{u} are interpreted as vectors in \mathbb{R}^3 . The normalized kernel is then

$$k_{3,\mathbf{s}}(\mathbf{u}) = \frac{\sqrt{3}}{\sqrt{4\pi\theta(\mathbf{s})^3}} \mathbb{I}\left(\|\mathbf{s} - \mathbf{u}\| < \frac{\theta(\mathbf{s})}{2}\right), \quad (4)$$

and the expression for the resulting taper is obtained as the normalized volume of the intersection of two spheres in \mathbb{R}^3 ,

$$T_3(\mathbf{s}, \mathbf{t}) = \frac{3}{4\pi(r_{st}R_{st})^{3/2}} \cdot \begin{cases} \frac{4\pi r_{st}^{3/2}}{3} & d_{\mathbf{s}\mathbf{t}} < R_{st} - r_{st}, \\ V_3\left(R_{st}, \frac{d_{st}^2 + R_{st}^2 - r_{st}^2}{2d_{st}}\right) + V_3\left(r_{st}, \frac{d_{st}^2 + r_{st}^2 - R_{st}^2}{2d_{st}}\right) & R_{st} - r_{st} \leq d_{\mathbf{s}\mathbf{t}} < R_{st} + r_{st}, \\ 0 & \text{otherwise.} \end{cases}$$

Here, $V_3(r, x)$ denotes the volume of the spherical cap with triangular height x of a sphere with radius r , which is given by

$$V_3(r, x) = \begin{cases} \frac{\pi}{3}(r-x)^2(2r+x) & |x| < r, \\ 0 & \text{otherwise.} \end{cases}$$

With $\theta(\mathbf{s}) \equiv \theta$, the taper simplifies to

$$T(d) = \begin{cases} \frac{1}{2\theta^3}(2\theta+d)(\theta-d)^2 & \text{if } d < \theta, \\ 0 & \text{otherwise,} \end{cases}$$

which is a valid covariance function on \mathbb{R}^m for $m \leq 3$ and is commonly referred to as the spherical covariance function.

By the same reasoning one can create a taper function by defining the generating kernel as a normalized hypersphere in \mathbb{R}^n . The taper is then obtained as the normalized volume of the intersection of the two hyperspheres, which for a general n is

$$T_n(\mathbf{s}, \mathbf{t}) = \frac{\Gamma(n/2+1)}{(\pi r_{st}R_{st})^{n/2}} \cdot \begin{cases} \frac{\pi^n r_{st}^{n/2}}{\Gamma(n/2+1)} & \text{if } d_{\mathbf{s}\mathbf{t}} < R_{st} - r_{st}, \\ V_n\left(R_{st}, \frac{d_{st}^2 + R_{st}^2 - r_{st}^2}{2d_{st}}\right) + V_n\left(r_{st}, \frac{d_{st}^2 + r_{st}^2 - R_{st}^2}{2d_{st}}\right) & \text{if } R_{st} - r_{st} \leq d_{\mathbf{s}\mathbf{t}} < R_{st} + r_{st}. \\ 0 & \text{otherwise.} \end{cases}$$

Here $\Gamma(x)$ is the gamma function and V_n denotes the volume of a hyperspherical cap with triangular height x , of a hypersphere with radius r , which is given by

$$V_n(r, x) = \frac{\pi^{n/2} r^n}{2\Gamma(n/2+1)} \cdot \begin{cases} I_{1-(x/r)^2}(\frac{n+1}{2}, \frac{1}{2}) & 0 \leq x < r, \\ 1 - I_{1-(x/r)^2}(\frac{n+1}{2}, \frac{1}{2}) & -r < x < 0, \\ 0 & \text{otherwise.} \end{cases}$$

Here, I_x is the regularized incomplete beta function, defined as

$$I_x\left(\frac{n+1}{2}, \frac{1}{2}\right) = \frac{\int_0^x \sqrt{t^{n-1}(1-t)^{-1}} dt}{\int_0^1 \sqrt{t^{n-1}(1-t)^{-1}} dt}.$$

See Li (2011) for a derivation of the volume of a hyperspherical cap.

By construction, $T_n(\mathbf{s}, \mathbf{t})$ is a valid covariance function for \mathbb{R}^m , $m \leq n$. Choosing $\theta(\mathbf{s}) \equiv \theta$ results in the stationary and isotropic correlation function

$$T_n(d) = I_n\left(1 - \left(\frac{d}{\theta}\right)^2\right), \quad (5)$$

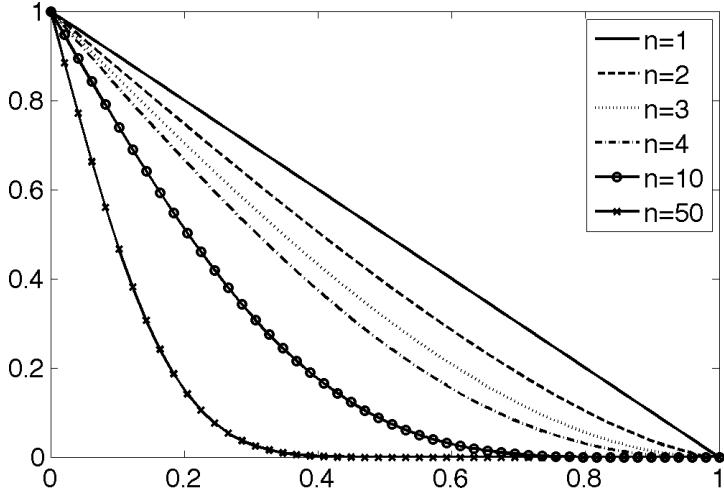


Figure 2: The stationary hyperspherical covariance functions for $\theta = 1$ for different values of n .

which is shown for different values of n in Figure 2. This function is sometimes referred to as Euclid's hat and is frequently used in spatial statistics (see e.g. Matérn, 1960, Chapter 3). The function is continuous at the origin and has $n - 1$ derivatives at $d = \theta$ (Gneiting, 1999). Furthermore, Euclid's hat with $n = m$ is the solution to Turán's problem, which means that it is the correlation function with the maximal integral over its support (see Ehm et al., 2004, Section 4.4). This is important for tapering because one then wants to taper low values to zero while keeping large value relatively unaffected. Thus, these tapers are both simple to construct and have several desirable properties.

2.2 Smooth tapers

The differentiability at the origin of the taper function is an important property. When using a compactly supported covariance function for tapering, one generally wants it to have as many derivatives at the origin as the original covariance function (Furrer et al., 2006). A drawback with the family T_n is therefore that all functions in the class have the same differentiability at the origin.

To obtain a smoother taper one can replace the spherical kernel with some differentiable function in (1). For example, one could define an m times differentiable kernel as the convolution of m spherical kernels, $k_{n,s}^m(\mathbf{u}) = k_{n,s}(\mathbf{u}) * \dots * k_{n,s}(\mathbf{u})$. However, the problem with this construction, and most other convolution-based approaches, is that the integral in (1) in general is not available in closed form. We believe that one could compute the taper analytically if $k_{n,s}^m(\mathbf{u})$ is used as a kernel, but it is not clear whether one can find some compact and simple expression for it.

An alternative is to do the integration numerically. However, as soon as numerical integration is required the computational cost for forming the taper matrix increases and can easily outweigh the cost of, for example, computing the kriging predictor using a full covariance matrix. We have not been able to find any non-stationary tapers based on isotropic kernels which are as easy to evaluate as the T_n family, and will not pursue this issue further here. See for example Mateu et al. (2013) for other kernel-based approaches to constructing non-stationary compactly supported covariance functions.

A simpler alternative for constructing a more differentiable non-stationary covariance function is to drop the requirement that the kernel should be isotropic. One can then construct the kernel as a product $k_s^m(\mathbf{u}) = \prod_{i=1}^d k_{s_i}^m(u_i)$, where $k_{s_i}^m(u) = k_{s_i}(u) * \dots * k_{s_i}(u)$ is the convolution

of m spherical kernels on \mathbb{R} . The taper is then given by

$$T^m(\mathbf{s}, \mathbf{t}) = \int k_{\mathbf{s}}^m(\mathbf{u}) k_{\mathbf{t}}^m(\mathbf{u}) d\mathbf{u} = \prod_{i=1}^d \int k_{s_i}^m(u_i) k_{t_i}^m(u_i) du_i = \prod_{i=1}^d T^m(s_i, t_i). \quad (6)$$

The spherical kernel on \mathbb{R} is simply an indicator function and k_s^m is therefore a scaled and shifted cardinal B-spline. Thus, to evaluate $T^m(s_i, t_i)$ one has to compute the integral of the product of two B splines, which can be done using integration by parts (Vermeulen et al., 1992).

For $m = 1$ one has $k_s^1(u) = \mathbb{I}(|u - s| < \theta(s)/2)$ and

$$T^1(s, t) = \frac{1}{\sqrt{\theta(s)\theta(t)}} \begin{cases} |s - t|, & \text{if } 2|s - t| < \theta(s) + \theta(t) \\ 0 & \text{otherwise.} \end{cases}$$

This is the one-dimensional version of the Hyperspherical taper $T_1(s)$, and if $\theta(s) \equiv \theta$ it is the triangular covariance function. For $m = 2$, the kernel instead is

$$k_s^2(u) = \frac{\sqrt{3}}{\sqrt{\theta(s)}} \begin{cases} 1 + 2\frac{u}{\theta(s)} & \text{if } s - \theta(s) < u < s, \\ 1 - 2\frac{u}{\theta(s)} & \text{if } s \leq u < s + \theta(s), \\ 0 & \text{otherwise.} \end{cases}$$

which gives the taper

$$T^2(s, t) = \frac{1}{4(rR)^{3/2}} \cdot \begin{cases} 2d^3 - 6d^2r + 2r^2(3R - r) & d \leq \min(r, R - r), \\ 3[d^3 - (d^2 + rR)(R + r) + d(r - R)^2] - R^3 - r^3 & R - r < d < r, \\ 6r^2(R - d) & r < d \leq R - r, \\ 3[(rR - d^2)(r + R) + d(R - r)^2] + r^3 - R^3 + d^3 & \max(r, R - r) < d \leq R, \\ 3[(d^2 + rR)(r + R) - d(r + R)^2] + R^3 + r^3 - d^3 & R < d \leq R + r, \\ 0 & \text{otherwise.} \end{cases}$$

Here, $d = |s - t|$, $r = \frac{1}{2} \min(\theta(s), \theta(t))$, and $R = \frac{1}{2} \max(\theta(s), \theta(t))$. If $\theta(s) \equiv \theta$, $T^2(s, t)$ simplifies to

$$T^2(d) = \frac{1}{8\theta^3} \cdot \begin{cases} \theta^3 + 6d^3 - 6\theta d^2 & 0 < d < \theta/2, \\ 2(\theta^3 - 3\theta^2 d + 3\theta d^2 - d^3) & \theta/2 < d \leq \theta, \\ 0 & \text{otherwise.} \end{cases}$$

3 Tapering and kriging prediction

Let $X(\mathbf{s})$ be a mean-zero Gaussian random field with some covariance function $r(\mathbf{s}, \mathbf{t})$. Suppose that X is observed at the locations $\mathbf{s}_1, \dots, \mathbf{s}_N$ and let $\mathbf{x}_0 = (X(\mathbf{s}_1), \dots, X(\mathbf{s}_N))^{\top}$. Further, let $\mathbf{x}_1 = (X(\mathbf{s}_1^*), \dots, X(\mathbf{s}_n^*))^{\top}$ be a vector with values that are to be predicted. The joint distribution of \mathbf{x}_0 and \mathbf{x}_1 can then be written as

$$\begin{pmatrix} \mathbf{x}_0 \\ \mathbf{x}_1 \end{pmatrix} \sim \mathcal{N} \left(\mathbf{0}, \begin{pmatrix} \Sigma_{00} & \Sigma_{01} \\ \Sigma_{10} & \Sigma_{11} \end{pmatrix} \right).$$

The conditional distribution of \mathbf{x}_1 given \mathbf{x}_0 is $\mathbf{x}_1 | \mathbf{x}_0 \sim \mathcal{N}(\Sigma_{10}\Sigma_{00}^{-1}\mathbf{x}_0, \Sigma_{11} - \Sigma_{10}\Sigma_{00}^{-1}\Sigma_{01})$. The mean $\hat{\mathbf{x}}_1 = \Sigma_{10}\Sigma_{00}^{-1}\mathbf{x}_0$ is the kriging predictor, and the MSE of the predictor is given by the diagonal of the covariance matrix, $\text{MSE}(\hat{\mathbf{x}}) = \text{diag}(\Sigma_{11} - \Sigma_{10}\Sigma_{00}^{-1}\Sigma_{01})$. As previously mentioned, solving $\Sigma_{00}^{-1}\mathbf{x}_0$ is computationally expensive, and an alternative is to use a tapering estimate $\tilde{\mathbf{x}}_1 = \tilde{\Sigma}_{10}\tilde{\Sigma}_{00}^{-1}\mathbf{x}_0$, where the tilde indicates that the covariance matrices are based on

the tapered covariance $r(\mathbf{s}, \mathbf{t})T(\mathbf{s}, \mathbf{t})$. That is, $\tilde{\Sigma}_{00} = \Sigma_{00} \circ \mathbf{T}$, where $T_{ij} = T(\mathbf{s}_i, \mathbf{s}_j)$ and \circ denotes the element-wise product. The MSE of the tapered prediction is given by

$$\text{MSE}(\tilde{\mathbf{x}}) = \text{diag}(\Sigma_{11} - \tilde{\Sigma}_{10}\tilde{\Sigma}_{00}^{-1}\Sigma_{01} - \Sigma_{10}\tilde{\Sigma}_{00}^{-1}\tilde{\Sigma}_{01} + \tilde{\Sigma}_{10}\tilde{\Sigma}_{00}^{-1}\Sigma_{00}\tilde{\Sigma}_{00}^{-1}\tilde{\Sigma}_{01}), \quad (7)$$

which reduces to the MSE of the kriging predictor if $T(\mathbf{s}, \mathbf{t}) = 1$. The MSE of the tapered prediction is a natural measure of accuracy which we will use to compare different methods in what follows.

3.1 Using adaptive tapers

When tapering is used for kriging, the stationary taper range is usually set so that $\tilde{\Sigma}_{00}$ has a fixed percentage of non-zero values. The goal with adaptive tapering is to achieve better predictions without increasing the number of non-zero values in $\tilde{\Sigma}_{00}$. To do this, the local taper ranges $\theta(\mathbf{s})$ of the adaptive taper have to be chosen. The best choice of $\theta(\mathbf{s})$ would be a function that minimizes the kriging error while keeping the number of non-zero elements in $\tilde{\Sigma}_{00}$ fixed. This is however difficult to do in practice, so we will instead split the problem in two parts. Since the sparsity of $\tilde{\Sigma}_{00}$ only is affected by the values of $\theta(\mathbf{s})$ at the measurement locations $\mathbf{s}_1, \dots, \mathbf{s}_N$, we start by choosing $\boldsymbol{\theta}_{1:N} = (\theta(\mathbf{s}_1), \theta(\mathbf{s}_2), \dots, \theta(\mathbf{s}_N))$ in order to achieve a given sparsity. As a second step, we interpolate the values $\boldsymbol{\theta}_{1:N}$ to obtain the continuous function $\theta(\mathbf{s})$ which is used to calculate $\tilde{\Sigma}_{10}$. The details of these two steps are discussed in the following two subsections.

3.1.1 Choosing the taper range at the measurement locations

Ideally, $\boldsymbol{\theta}_{1:N}$ should be chosen so that every row in $\tilde{\Sigma}_{00}$ has the same number, m , of non-zero elements. In other words, one would like the tapered covariance between $X(\mathbf{s}_i)$ and X at the $m-1$ points nearest to \mathbf{s}_i to be non-zero for each i . In general, one cannot get exactly m elements for each row since the nearest-neighbor relation is not symmetric. However, $\boldsymbol{\theta}_{1:N}$ can be chosen so that the average number per row is m while keeping the variation between the rows small. This can be achieved using Algorithm 1, where $n_j(\boldsymbol{\theta}_{1:N})$ is the number of non-zero elements of row j in $\tilde{\Sigma}_{00}$ for a given choice of taper ranges $\boldsymbol{\theta}_{1:N}$, and $r_j(\boldsymbol{\theta}_{1:N}, k)$ denotes the k th largest value of the vector $(\|\mathbf{s}_j - \mathbf{s}_1\|, \|\mathbf{s}_j - \mathbf{s}_2\|, \dots, \|\mathbf{s}_j - \mathbf{s}_N\|)^T - \boldsymbol{\theta}_{1:N}$. The parameter $\mathbf{m}_{1:N} = (m_1, \dots, m_N)$ is a vector containing the desired number of non-zero elements for each row. The parameter ϵ determines the percentage of nodes that is allowed to have $|n_j(\boldsymbol{\theta}_{1:N}) - m_j| > 1$, and thus serves as a convergence criterium. Unless the measurement locations are structured in some malicious way, the algorithm will converge towards a point $\boldsymbol{\theta}_{1:N}$ that gives close to m_j non-zero elements for each row j .

The algorithm works as follows: If $\tilde{\Sigma}_{00}$ has too many non-zero elements at iteration i , the taper range of the location which corresponds to the row with the most non-zero elements is decreased so that that row gets the correct number of non-zero elements. If the matrix instead has too few non-zero elements at iteration i , the taper range of the location which corresponds to the row with the fewest number of non-zero elements is increased so that that row gets the correct number of non-zero elements. In the case that $\tilde{\Sigma}_{00}$ has the correct number of non-zero elements, but the rows have different number of non-zero elements, the row with the largest deviation in the number of non-zero elements is changed. If several rows have the same deviation, one of these are picked at random at row 16. In row 17, random numbers are used since deterministic values can cause taper ranges for different points to be exactly equal, which can cause numerical problems. When the first loop has converged, the final loop goes through the nodes and maximizes the taper ranges while keeping the number of non-zero elements constant.

Algorithm 1

```

1: procedure THETA-SET( $\mathbf{m}_{1:N}, \boldsymbol{\theta}_{1:N}^{(0)}, \mathbf{s}_{1:N}, \epsilon$ )
2:    $i \leftarrow 0$ 
3:   do
4:      $i \leftarrow i + 1$ 
5:      $\boldsymbol{\theta}_{1:N}^{(i)} \leftarrow \boldsymbol{\theta}_{1:N}^{(i-1)}$ 
6:      $n_m \leftarrow \min_l(n_l(\boldsymbol{\theta}_{1:N}^{(i)}) - m_l)$ 
7:     if  $\sum_l n_l(\boldsymbol{\theta}_{1:N}^{(j-1)}) = M$  then
8:       if  $\sum_l |n_l(\boldsymbol{\theta}_{1:N}^{(i)}) - m_l| = 0$  then
9:         break
10:        else if  $\max_l(n_l(\boldsymbol{\theta}_{1:N}^{(i)}) - m_l) > |\min_l(n_l(\boldsymbol{\theta}_{1:N}^{(i)}) - m_l)|$  then
11:           $n_m \leftarrow \max_l(n_l(\boldsymbol{\theta}_{1:N}^{(i)}) - m_l)$ 
12:        end if
13:        else if  $\sum_l n_l(\boldsymbol{\theta}_{1:N}^{(i)}) > Nm$  then
14:           $n_m \leftarrow \max_l(n_l(\boldsymbol{\theta}_{1:N}^{(i)}) - m_l)$ 
15:        end if
16:         $j \leftarrow$  random integer from the set  $\{l : n_l(\boldsymbol{\theta}_{1:N}^{(i)}) - m_l = n_m\}$ 
17:         $\theta_j \leftarrow \text{Unif}\left(r_j(\boldsymbol{\theta}_{1:N}^{(i)}, [m_j]), r_j(\boldsymbol{\theta}_{1:N}^{(i)}, [m_j] + 1)\right)$ 
18:      while  $\{\sum_l \mathbb{I}(|n_l(\boldsymbol{\theta}_{1:N}^{(i)}) - m_j| > 1) > N\epsilon\}$ 
19:      for  $j = 1, \dots, N$  do
20:         $\theta_j \leftarrow r_j(\boldsymbol{\theta}_{1:N}^{(i)}, n_j(\boldsymbol{\theta}_{1:N}^{(i)}))$ 
21:      end for
22:      return  $\boldsymbol{\theta}_{1:N}^{(i)}$ 
23: end procedure

```

The result of the algorithm for a simple example with ten locations can be seen in the left panel of Figure 3. The result is visualized through the circular kernels $k_{\mathbf{s}}(\mathbf{u})$ for each location, where the radius of the circle at location \mathbf{s}_i is given by $\theta_{\mathbf{s}_i}/2$.

3.1.2 Interpolating the taper ranges

Running Algorithm 1 results in a vector $\boldsymbol{\theta}_{1:N}$ that has to be interpolated in order to generate a continuous function $\theta(\mathbf{s})$. There are many ways to do this, and one of the simplest is to compute a Delaunay triangulation of the locations $\mathbf{s}_1, \dots, \mathbf{s}_N$ and do linear interpolation within each triangle: For a location \mathbf{s} within a triangle \mathcal{T} with vertices $\mathbf{s}_a, \mathbf{s}_b, \mathbf{s}_c$ and corresponding taper ranges $\theta_a, \theta_b, \theta_c$, the interpolated taper range $\theta(\mathbf{s})$ is given by $\theta(\mathbf{s}) = w_a\theta_a + w_b\theta_b + w_c\theta_c$, where $\{w_a, w_b, w_c\}$ are the barycentric coordinates of \mathbf{s} in \mathcal{T} . The result of this interpolation method, implemented using the `scatteredInterpolant` function in MATLAB (2015), for a simple example with ten locations can be seen in the right panel of Figure 3.

One could imagine that more advanced interpolation methods, resulting in smoother surfaces, could improve the results when the method is used for kriging. However, initial studies indicate that the interpolation method has little effect on the resulting Kriging error. For example, the `scatteredInterpolant` function also supports a natural neighbor interpolation method that uses Voronoi tessellations to find a smoother interpolation (C^1 except at the measurement locations). In order to investigate if the kriging results could be improved by using more sophisticated interpolation methods, a part of the simulation study in Section 3.2 was done using both natural and linear interpolation. Specifically, the MSE of the kriging predictor using linear interpolation for one of the test cases using an exponential covariance function was

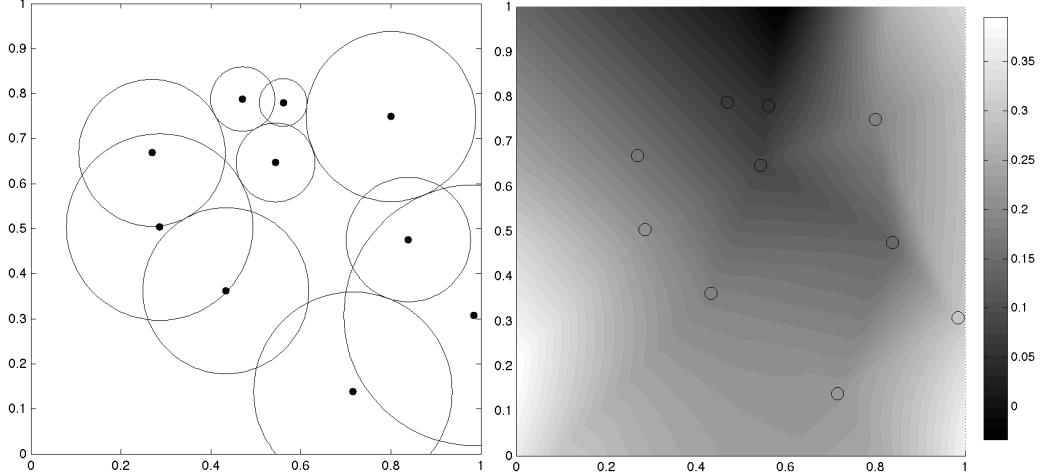


Figure 3: The left panel shows ten observation locations and the circular kernels $k_s(\mathbf{u})$ for each observation location. The diameters $\theta(s)$ of the kernels are chosen using Algorithm 1 so that each kernel overlaps with two other kernels, resulting in a tapering matrix with three non-zero elements per row. The right panel shows the taper range of the observation locations (the values in the circles) and an interpolated function which is used during kriging predictions of unobserved locations.

2.87 (shown as 2.9 in Table 1). Changing to the natural interpolation method instead resulted in an MSE of 2.78. The mean time it took to compute the interpolation was $9.1ms$ for the linear interpolation and $12.6ms$ for the natural interpolation. Thus, it seems as if the results for the adaptive tapers could be improved slightly by using more complicated interpolation methods. However, since the difference is small, we use the simpler linear interpolation method for the rest of this work, and leave the question of how to optimally interpolate the taper ranges for future research.

3.1.3 Examples using simulated data

Results of using Algorithm 1 in combination with linear interpolation for three realistic scenarios are shown in Figure 4. The left panels show the measurement locations and the right panels show the corresponding taper ranges chosen so that the tapered covariance matrix has as many non-zero elements as if a stationary taper range of 0.1 was used.

The top row shows a spatially structured case, where the measurements are taken at a perturbed grid. The measurements are done at the locations $\frac{1}{32}(0.5 + i + U_{ij}, 0.5 + j + V_{ij})$ for $i, j = 0, \dots, 31$, where U_{ij} and V_{ij} are uniform random variables on $(-0.45, 0.45)$. Thus, this scenario is similar to that used in Stein (2013). For this scenario, the taper range is larger at the edges of the domain but fairly constant in the interior, and the results are fairly similar to the warping method proposed by Stein (2013). The second scenario, shown in the middle row of Figure 4, is complete spatial randomness, where the measurements are taken at (U_{ij}, V_{ij}) where U_{ij} and V_{ij} now are uniform random variables on $(0, 1)$. Finally, the third scenario shown in the bottom row in the figure, corresponds to clustered locations, where the measurements locations are drawn from a log-Gaussian cox process. That is, the locations are drawn from a Poisson process with intensity $\lambda(s) = \exp(Z(s))$ where $Z(s)$ is a mean-zero Gaussian Matérn field, with covariance function

$$C(\mathbf{h}) = \frac{\sigma^2 2^{1-\nu}}{\Gamma(\nu)} (\kappa \|\mathbf{h}\|)^\nu K_\nu(\kappa \|\mathbf{h}\|), \quad \mathbf{h} \in \mathbb{R}^d, \nu > 0. \quad (8)$$

Here κ is a range parameter, σ^2 the variance, ν is a smoothness parameter, and K_ν is a modified Bessel function. The distance where the covariance is approximately one tenth of the variance, sometimes referred to as the practical range, of the field is given by $\rho = \sqrt{8\nu}\kappa^{-1}$. For the scenario in Figure 4, the parameter $\kappa = 10$, $\nu = 0.5$, and $\sigma = 2$ were used.

Compared with the spatially structured case one can note larger variations in the taper ranges in the two other scenarios. The warping method in Stein (2013) cannot be used to construct tapers similar to our adaptive tapers in these cases, but one would then have to use the more complicated and computationally expensive warping methods by Anderes et al. (2013).

Histograms of the number of non-zero elements per row can be seen in Figure 5. The figure has one panel for each of the scenarios in Figure 4, and each panel shows the histogram for the adaptive taper in black and the histogram for the stationary taper in grey. The total number of non-zero elements is the same for the two methods, but the variability of the number of non-zero elements for each row is much smaller for the adaptive tapers, especially in the clustered scenario. The variability of the number of non-zero elements per row for the adaptive taper is mainly caused by the fact that the average number of non-zero elements per row is not an integer. For example, for the leftmost panel it was approximately 29.6, and in order to get the correct average the method needs to let the rows have different number of non-zero elements.

3.2 Numerical comparisons of kriging prediction

In this section, stationary and adaptive tapering approximations are compared in the case of kriging prediction for a Gaussian Matérn field on $[0, 1]^2$ with known parameters. The field is observed at N locations in the domain and the value of the field is predicted for all locations on a 50×50 regular lattice in the domain. The comparisons are done both with $\nu = 0.5$, which corresponds to an exponential covariance function, and with $\nu = 1.5$ which corresponds to a smoother covariance function. For both values of ν , one case with a long correlation range, $\rho = 0.2$, and one case with a shorter range, $\rho = 0.1$, are tested. For all cases, $\sigma = 1$ is used. For each parameter configuration, the three different measurement scenarios described in Section 3.1.3 are tested, each with $N = 1024$ measurement locations.

Five different tapering functions are compared. The first is a stationary Wendland function

$$r_\theta(h) = \left(1 - \frac{|h|}{\theta}\right)^4 \left(1 + 4\frac{|h|}{\theta}\right) \mathbb{I}(|h| < \theta),$$

the second is a stationary hyperspherical taper (5) with $n = 2$, and the third is the corresponding non-stationary hyperspherical taper T_2 . The final two are the non-stationary product tapers T^1 and T^2 , where the latter has the same differentiability at the origin as the Wendland function. In order to get comparable results, the various taper ranges have to be chosen so that the methods have similar computational costs. To do this, the taper ranges of the two stationary tapers are set to 0.1. Algorithm 1 is then used to choose the non-stationary taper ranges so that all tapered covariance matrices have the same number of non-zero elements.

In summary, four different parameter settings and three different measurement scenarios are tested. This gives twelve test cases in total on which all five tapering methods are tested. For each test case, 100 different data sets are simulated, and the optimal kriging predictor and the different tapered approximations are computed for each data set.

The relative increase in the MSE, compared to the optimal kriging predictor, is used as a measure of accuracy

$$\frac{1}{100n} \sum_{j=1}^{100} \sum_{i=1}^n \frac{\text{MSE}(\tilde{\mathbf{x}}_0^{(j)}(\mathbf{s}_i)) - \text{MSE}(\hat{\mathbf{x}}_0^{(j)}(\mathbf{s}_i))}{\text{MSE}(\hat{\mathbf{x}}_0^{(j)}(\mathbf{s}_i))}, \quad (9)$$

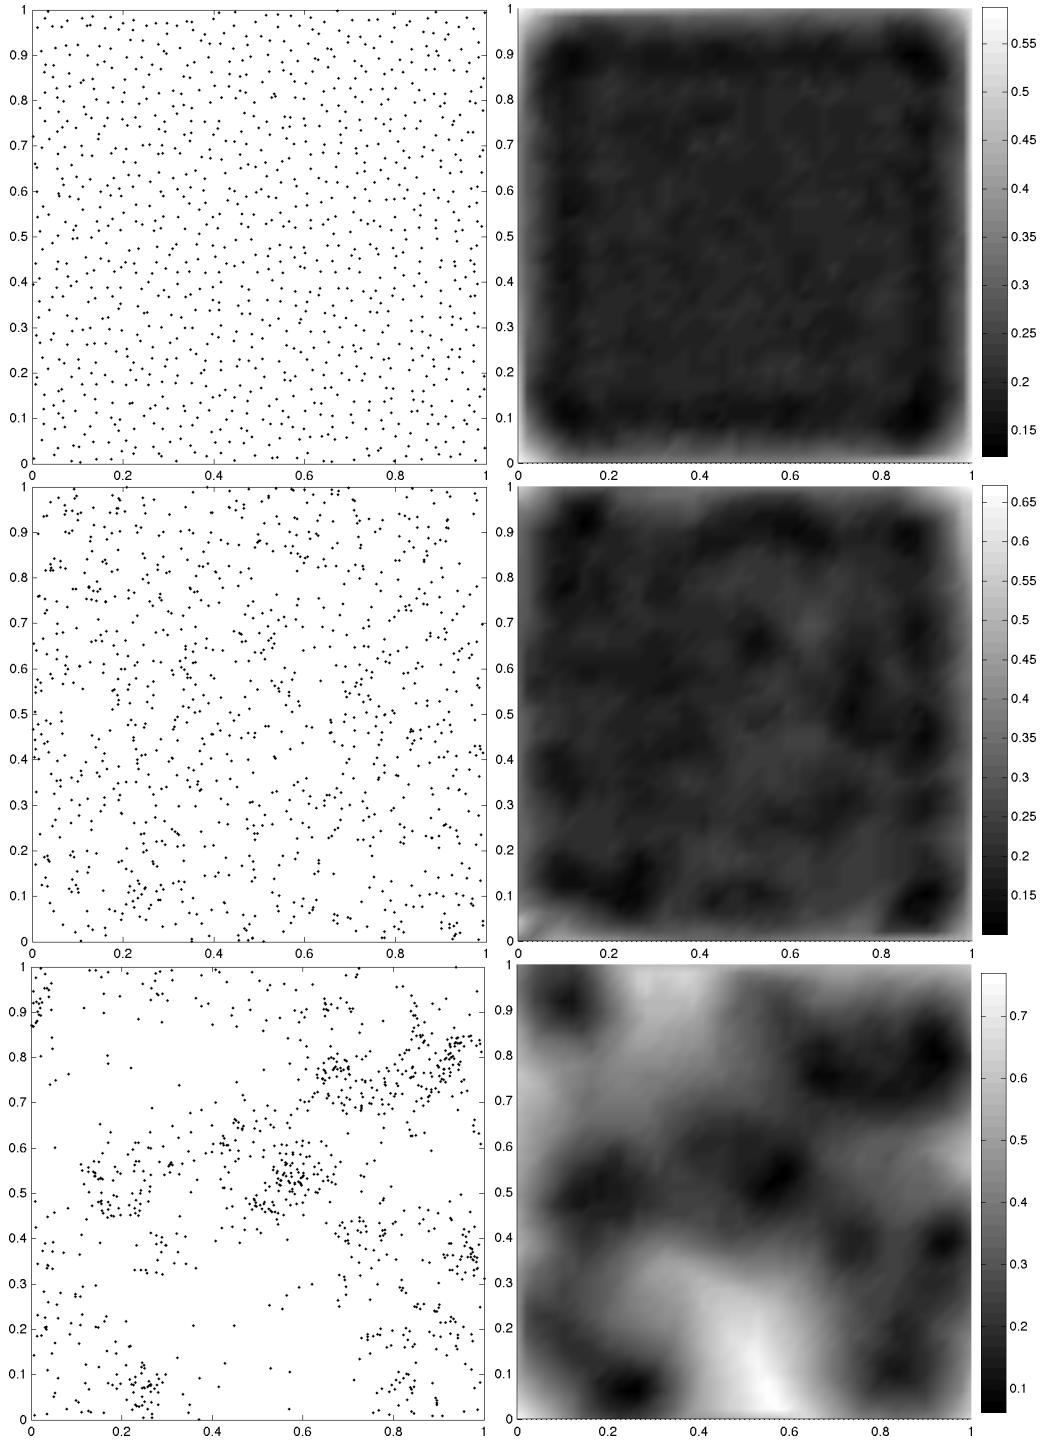


Figure 4: The left panels show three types of measurement locations: Structured locations (top), completely spatially random locations (mid), and clustered locations (bottom). Each example has 1024 measurement locations and the right panels show the corresponding non-stationary taper range for each scenario. The taper range is chosen using the method in Section 3.1 so that the covariance matrix for the data has as many non-zero elements as if a stationary taper range of 0.1 was used.

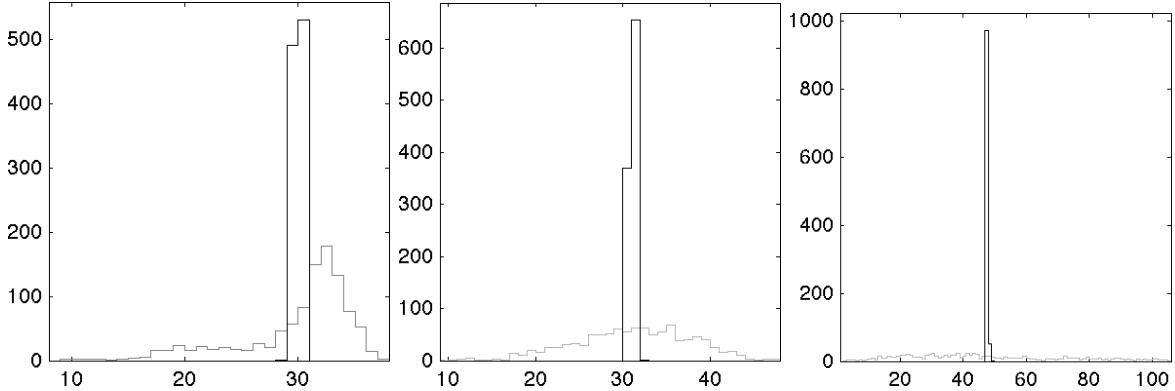


Figure 5: Histograms of the number of non-zero elements per row in the tapered covariance matrix for data for the three examples in Figure 4: Structured locations (left), random locations (middle), and clustered locations (right). The grey curves show the stationary tapers and the black curves show the adaptive tapers.

	Short range			Long range		
	S	R	C	S	R	C
$W(h)$	3.1 (0.06)	5.9 (0.18)	9.3 (0.42)	6.3 (0.11)	12.3 (0.52)	26.6 (2.15)
$T_2(h)$	1.2 (0.02)	2.3 (0.08)	4.7 (0.34)	1.8 (0.05)	3.7 (0.22)	12.5 (1.64)
$T_2(s, t)$	1.1 (0.02)	2.0 (0.08)	1.2 (0.33)	1.5 (0.03)	2.9 (0.15)	1.8 (0.58)
$T^1(s, t)$	1.6 (0.02)	2.5 (0.09)	1.5 (0.37)	2.8 (0.04)	4.0 (0.15)	2.6 (0.65)
$T^2(s, t)$	3.3 (0.06)	5.7 (0.22)	2.6 (0.88)	7.3 (0.11)	12.2 (0.50)	5.8 (2.12)

Table 1: The relative increase in MSE shown in percent (equation (9) scaled by a factor 100), for the different tapered kriging predictors compared to the optimal kriging predictor for the exponential covariance function. The results are shown for the Wendland taper (W), the stationary hyperspherical taper $T_2(h)$, the non-stationary hyperspherical taper $T_2(s, t)$, and the two product tapers $T^1(s, t)$ and $T^2(s, t)$ for structured (S), random (R), and clustered (C) sampling scenarios. The values in the parentheses are the Monte Carlo standard deviation for each estimate. The bold values indicate the best method for each case.

where $\tilde{\mathbf{x}}_0^{(j)}(\mathbf{s}_i)$ is the tapered kriging prediction at location \mathbf{s}_i for the simulated dataset j . These values are shown in Table 1 for the exponential covariance function, and in Table 2 for the Matérn covariance function.

There are several things to note in the tables. For the exponential covariance the non-stationary Hyperspherical taper is always better than the stationary tapers. The improvement is relatively small for structured observation locations, where it mainly comes from better predictions near the boundary of the observation domain, but the improvement is larger for spatially random locations and the largest for clustered observation locations. One should generally choose a tapering function which has the same differentiability as the true covariance function (see e.g. Furrer et al., 2006), and this is the reason for why the hyperspherical tapers outperform the Wendland taper and the product taper T^2 in the exponential case. Also for the Matérn covariance, the adaptive tapers perform the best in general. In this case, however, the product taper T^2 outperforms the hyperspherical taper. The reason for this is that the hyperspherical tapers do not satisfy the tapering condition by Furrer et al. (2006) for this case, and they therefore do not have asymptotically equivalent MSEs to the optimal kriging predictor. However, it is interesting to note that the non-stationary hyperspherical taper outperforms the Wendland taper in the clustered situations, despite the fact that it does not satisfy the tapering

	Short range			Long range		
	S	R	C	S	R	C
$W(h)$	6.7 (0.22)	16.0 (0.58)	23.7 (1.23)	22.5 (0.90)	60.5 (3.43)	131.7 (8.25)
$T_2(h)$	14.8 (0.20)	21.6 (0.50)	23.7 (0.81)	46.7 (0.60)	62.4 (1.52)	89.9 (4.28)
$T_2(s, t)$	15.0 (0.23)	21.9 (0.56)	14.7 (2.65)	45.1 (0.56)	59.9 (1.62)	42.3 (6.59)
$T^1(s, t)$	20.0 (0.26)	26.2 (0.58)	17.5 (2.96)	61.4 (0.73)	73.6 (1.81)	52.1 (7.12)
$T^2(s, t)$	6.5 (0.24)	15.3 (0.82)	6.3 (2.34)	26.8 (1.18)	64.1 (4.31)	28.3 (10.61)

Table 2: The relative increase in MSE shown in percent (equation (9) scaled by a factor 100), for the different tapered kriging predictors compared to the optimal kriging predictor for the Matérn covariance function. The bold values indicate the best method for each case. See Table 1 for explanations of the different test cases.

	Compute $\theta(\mathbf{s})$	Construct \mathbf{T}	Kriging
$W(h)$	—	445	32
$T_2(h)$	—	741	8
$T_2(s, t)$	228	797	11
$T^1(s, t)$	228	642	28
$T^2(s, t)$	228	817	27

Table 3: Average timings in milliseconds for the different steps in the simulation study. The first column shows the time for computing the non-stationary taper range using Algorithm 1 and interpolation. The second column shows the time for constructing the tapering matrix, and the third column shows the time for computing the kriging predictor using the tapered covariance matrix.

condition. It is also interesting to note that T^2 is best overall for the Matérn case, despite the fact that it is anisotropic. Thus, non-stationarity of the tapers seems to be important for the accuracy of the predictions.

Table 3 shows timing results of the different methods in the simulation study, implemented in MATLAB (2015) and performed using a Macbook Pro computer with a 2.6GHz Intel Core i7 processor (Apple Inc., Cupertino, CA, USA). The average time for computing $\theta(\mathbf{s})$ was 228ms, of which approximately 219ms was spent on running Algorithm 1, and the rest was used for interpolating the result. However, the timing for this step is dependent on which measurement scenario that is used, and the structured scenario resulted in the fastest average computation time (128ms) while the clustered was the slowest (376ms). The time for constructing the covariance matrices, given the taper range, was very similar for the different methods, and comparable to the time it took to run Algorithm 1. Also the time it took to compute the kriging predictor was similar for the different methods, and the reason for this is that all tapering matrices had the same number of non-zero elements. It should be stressed here that the timings for the first two steps are highly implementation dependent, and especially the time it takes to run Algorithm 1 could be greatly reduced by implementing the procedure in C or Fortran.

4 Tapering and parameter estimation

In the previous section, tapering was used for spatial prediction based on measurements of a Gaussian random field with known covariance function. In practice the parameters, Ψ , of the covariance function often have to be estimated from the data first, and we now look at how tapering can be used in this step. As before, \mathbf{x}_0 denotes a vector of observations of a mean-zero

Gaussian random field $X(\mathbf{s})$. The log-likelihood, that should be maximized to estimate the parameters, is $l(\Psi; \mathbf{x}_0) = -\frac{1}{2} \log |\Sigma_{00}| - \frac{1}{2} \mathbf{x}_0^\top \Sigma_{00}^{-1} \mathbf{x}_0$. The computational cost for evaluating the likelihood is $\mathcal{O}(N^3)$ due to both the log-determinant and the matrix inverse. The natural tapering approximation of the likelihood is to replace Σ_{00} with the tapered covariance matrix $\tilde{\Sigma}_{00}$. However, this can cause substantial bias in the resulting estimates (Kaufman et al., 2008), so a common choice is to instead use the tapered likelihood $\tilde{l}(\Psi; \mathbf{x}_0) = -\frac{1}{2} \log |\tilde{\Sigma}_{00}| - \frac{1}{2} \mathbf{x}_0^\top (\tilde{\Sigma}_{00}^{-1} \circ \mathbf{T}) \mathbf{x}_0$.

In the case of kriging estimation, adaptive tapers with taper ranges that adapted to the measurement locations was used to improve the predictions. For parameter estimation, however, it is not clear whether this is the best option. We therefore propose the following method for selecting tapering ranges in order to generate a taper matrix with M non-zero elements. The method, denoted α -adaptive tapering, can be used as a tradeoff between fully adaptive tapering and stationary tapering.

Let $\mathbf{m}_s = (m_1, \dots, m_N)$ be a vector where m_i is the number of non-zero elements for the i th row in \mathbf{T} for a stationary taper with taper range chosen so that \mathbf{T} has M non-zero elements in total, and select $\alpha \in [0, 1]$. Then, compute the adaptive taper ranges by running Algorithm 1 with $\mathbf{m} = (1 - \alpha)\mathbf{m}_s + \alpha M/N$ as the vector of desired elements per row. Using $\alpha = 0$ results in a stationary taper whereas $\alpha = 1$ results in the fully adaptive taper. Note that the method does not depend on the covariance matrix that is to be tapered, which means that the tapering matrix can be computed once before estimating the parameters.

4.1 Numerical comparisons for parameter estimation

In this section, adaptive tapering is tested in the case of parameter estimation for two exponential covariance models. Since exponential covariances are used, the hyperspherical taper T_2 is used in combination with the α -adaptive method for selecting the taper ranges. As a reference method, the observations are partitioned into blocks and the true log-likelihood is approximated by summing the log-likelihoods of the blocks. This can be implemented as a block-diagonal taper with elements $T_{ij} = 0$ if \mathbf{s}_i and \mathbf{s}_j are in the same block, and $T_{ij} = 0$ otherwise. Because of this, Stein (2013) refers to the method as block tapering.

The first test case is estimation of a stationary exponential covariance model, (8) with parameters $\nu = 0.5$, $\kappa = 10$, and $\sigma = 10$. The second test case is a non-stationary exponential covariance model. Specifically, the approach by Paciorek and Schervish (2006) is used to produce a non-stationary exponential covariance function

$$C(\mathbf{s}_i, \mathbf{s}_j) = \sigma^2 |\Sigma(\mathbf{s}_j)|^{1/4} |\Sigma(\mathbf{s}_j)|^{1/4} \left| \frac{\Sigma(\mathbf{s}_j) + \Sigma(\mathbf{s}_j)}{2} \right|^{-1/2} \exp \left(-\sqrt{Q_{ij}} \right). \quad (10)$$

Here σ^2 is the variance, which is set to 4 in the comparison, and

$$Q_{ij} = (\mathbf{s}_i - \mathbf{s}_j)^T \left(\frac{\Sigma(\mathbf{s}_j) + \Sigma(\mathbf{s}_j)}{2} \right)^{-1} (\mathbf{s}_i - \mathbf{s}_j).$$

For simplicity, we choose $\Sigma(\mathbf{s}) = \kappa(s) \mathbf{I}$ where $\kappa(s)$ is the spatially varying correlation range, modeled using a basis expansion $\log \kappa(\mathbf{s}) = \log \kappa(s_x, s_y) = \kappa_1 + \kappa_2 s_x$ where $\kappa_1 = -6$ and $\kappa_2 = 6$ are used. This results in a model where the covariance range is constant in the s_y direction but varies in the s_x direction.

For both the stationary and non-stationary test cases, data is generated by sampling a mean-zero Gaussian process $X(\mathbf{s})$ with the respective covariance function at $N = 1024$ locations on $[0, 1]^2$, chosen at random according to the spatially structured case used in Section 3.2. The number of blocks for the block taper and the taper ranges for the α -adaptive method are chosen so that all taper matrices have approximately $0.01N^2$ non-zero elements.

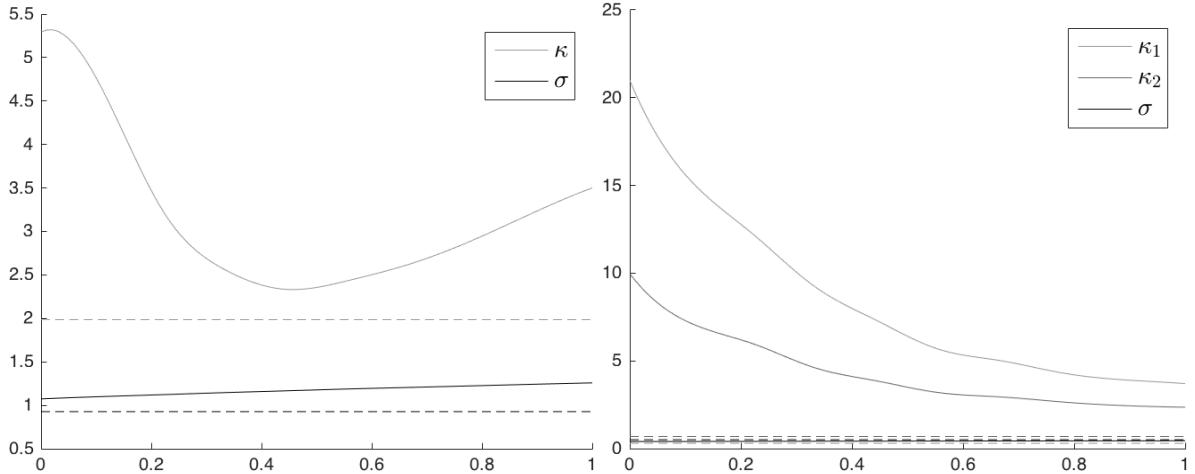


Figure 6: Tapering results for the stationary (left panel) and non-stationary (right panel) exponential models. Square roots of the diagonal elements of the inverse Godambe information matrix for the α -adaptive taper as a function of α (solid lines), and for the block taper (dashed lines). Note that ordinary stationary tapering corresponds to the case $\alpha = 0$.

We use the diagonal of the inverse Godambe information matrix (see e.g. Stein, 2013) as a measure of the efficiency of the estimation methods based on different tapering functions. The elements correspond to the asymptotic marginal variances of the estimates obtained using the tapered likelihoods. The Godambe information for the tapered loglikelihood is $\mathbf{G}(\Psi) = \mathbb{E}_\Psi(\mathbf{H})\mathbb{E}_\Psi(\mathbf{v}\mathbf{v}^\top)^{-1}\mathbb{E}_\Psi(\mathbf{H})$, where $\mathbf{v} = \nabla_\Psi \tilde{l}(\Psi; \mathbf{x}_0)$ and $\mathbf{H} = \Delta_\Psi \tilde{l}(\Psi; \mathbf{x}_0)$. The two expectations can be computed as $\mathbb{E}_\Psi(\mathbf{v}\mathbf{v}^\top)_{ij} = \frac{1}{2} \text{tr} \left((\Sigma_i^{-1} \circ \mathbf{T}) \Sigma_{00} (\Sigma_j^{-1} \circ \mathbf{T}) \Sigma_{00} \right)$ and $\mathbb{E}_\Psi(\mathbf{H})_{ij} = \frac{1}{2} \text{tr} \left(\Sigma_i^{-1} \Sigma_j \right)$, where $\Sigma_i^{-1} = -\tilde{\Sigma}_{00}^{-1} \Sigma_i \tilde{\Sigma}_{00}^{-1}$ and $\Sigma_i = \frac{\partial \tilde{\Sigma}_{00}}{\partial \Psi_i} = \frac{\partial \Sigma_{00}}{\partial \Psi_i} \circ \mathbf{T}$.

Figure 6 displays the square roots of the diagonal elements of the inverse Godambe information matrix, as functions of α for the α -adaptive method together with the corresponding values for the block taper for the two test cases. Recall that $\alpha = 0$ corresponds to stationary tapering whereas $\alpha = 1$ corresponds to fully adaptive tapering that minimizes the variation of non-zero elements between rows in the tapering matrix. The left panel of the figure shows the test case for the stationary exponential model, and one can note that the adaptive tapers outperform the stationary taper, and that an optimal value of α in this case is around 0.4. The right panel shows the non-stationary test case. The adaptive tapers outperform the stationary taper also in this test case, and the optimal value is in this case 1, meaning that the fully adaptive taper should be used. This is of little value, however, as the block taper is even better than the α -adaptive tapers for both cases.

These are of course just two examples and different choices of true covariance function and measurement locations would have resulted in different optimal values for α . However, the general conclusion has been the same for most cases we have tried so far: Non-stationary tapering outperforms stationary tapering, but the block tapers outperform both.

5 Discussion

We have presented a class of computationally convenient compactly supported non-stationary covariance functions. Together with the proposed methods for selecting non-stationary taper ranges, these functions can be used for spatially adaptive covariance tapering.

Numerical experiments showed that the adaptive tapering method can improve the performance of tapering for both kriging and parameter estimation, with negligible increase in com-

putational cost. For parameter estimation, however, the numerical experiments also showed that dividing the data into blocks and ignoring the dependence between the blocks is often a better method than tapering. Thus, even if adaptive tapering is better than stationary tapering, we agree with Stein (2013) that tapering seems to have little practical value for parameter estimation.

An alternative to using the adaptive tapers for parameter estimation in applications is to use them directly as non-stationary covariance models for the data. In this case, an interesting model extension is to allow for anisotropy in the covariance function by using elliptical kernels in the construction. With this straightforward extension, the model could be used as a computationally efficient alternative to the non-stationary models by Paciorek and Schervish (2006). Another relevant extension is to derive the expressions for the differentiable tapers, obtained using kernels that are convolutions of the spherical kernels. This would further increase the value of this class of covariance models. Finally, another interesting direction of future research is to use the adaptive tapers in an adaptive MCMC method similar to the method by Wallin and Bolin (2015).

Acknowledgements

The authors are grateful to the reviewers for their helpful and constructive comments on earlier versions of the manuscript. Both authors have been supported by the Knut and Alice Wallenberg foundation and the second author has been supported by the Swedish Research Council Grant 2008-5382.

References

- Anderes, E., Huser, R., Nychka, D., Coram, M., 2013. Nonstationary positive definite tapering on the plane. *Journal of Computational and Graphical Statistics* 22 (4), 848–865.
- Barry, R. P., Ver Hoef, J. M., 1996. Blackbox kriging: Spatial prediction without specifying variogram models. *J. Agr. Biol. Environ. Statist.* 1 (3), 297–322.
- Bolin, D., Lindgren, F., 2013. A comparison between Markov approximations and other methods for large spatial data sets. *Comput. Statist. and Data Anal.* 61, 7–21.
- Cressie, N., Ravlicová, M., 2002. Calibrated spatial moving average simulations. *Statist. Model.* 2, 267–279.
- Du, J., Zhang, H., Mandrekar, V., et al., 2009. Fixed-domain asymptotic properties of tapered maximum likelihood estimators. *The Annals of Statistics* 37 (6A), 3330–3361.
- Ehm, W., Gneiting, T., Richards, D., 2004. Convolution roots of radial positive definite functions with compact support. *Transactions of the American Mathematical Society* 356 (11), 4655–4685.
- Furrer, R., Genton, M. G., Nychka, D., 2006. Covariance tapering for interpolation of large spatial datasets. *J. Comput. Graph. Statist.* 15, 502–523.
- Gneiting, T., 1999. Radial positive definite functions generated by euclid's hat. *Journal of Multivariate Analysis* 69 (1), 88–119.
- Gneiting, T., 2002. Compactly supported correlation functions. *J. Multivariate Anal.* 83, 493–508.
- Higdon, D., 2001. Space and space-time modeling using process convolutions. Tech. Rep. 01-03, Duke University, Durham, NC.

- Kaufman, C. G., Schervish, M. J., Nychka, D. W., 2008. Covariance tapering for likelihood-based estimation in large spatial data sets. *Journal of the American Statistical Association* 103 (484), 1545–1555.
- Li, S., 2011. Concise formulas for the area and volume of a hyperspherical cap. *Asian Journal of Mathematics and Statistics* 4 (1), 66–70.
- Matérn, B., 1960. Spatial variation. *Meddelanden från statens skogsforskningsinstitut* 49 (5).
- Mateu, J., Fernández-Avilés, G., Montero, J., 2013. On a class of non-stationary, compactly supported spatial covariance functions. *Stochastic Environmental Research and Risk Assessment* 27 (2), 297–309.
- MATLAB, 2015. 8.6.0.267246 (R2015b). The MathWorks Inc., Natick, Massachusetts.
- Paciorek, C. J., Schervish, M. J., 2006. Spatial modelling using a new class of nonstationary covariance functions. *Environmetrics* 17 (5), 483–506.
- Rodrigues, A., Diggle, P. J., 2010. A class of convolution-based models for spatio-temporal processes with non-separable covariance structure. *Scand. J. Statist.* 37, 553–567.
- Shaby, B., Ruppert, D., 2012. Tapered covariance: Bayesian estimation and asymptotics. *Journal of Computational and Graphical Statistics* 21 (2), 433–452.
- Stein, M. L., 2013. Statistical properties of covariance tapers. *Journal of Computational and Graphical Statistics* 22 (4), 866–885.
- Vermeulen, A. H., Bartels, R. H., Heppler, G. R., 1992. Integrating products of b-splines. *SIAM journal on scientific and statistical computing* 13 (4), 1025–1038.
- Wallin, J., Bolin, D., 2015. Efficient adaptive MCMC through precision estimation. arXiv preprint arXiv:1505.03908.
- Wang, D., Loh, W.-L., et al., 2011. On fixed-domain asymptotics and covariance tapering in gaussian random field models. *Electronic Journal of Statistics* 5, 238–269.
- Wendland, H., 1995. Piecewise polynomial, positive definite and compactly supported radial functions of minimal degree. *Adv. Comput. Math.* 4, 389–396.